AD-A206 864

THE CONTRIBUTIONS OF J H WILKINSON TO NUMERICAL
ANALYSIS(U) CALIFORNIA UNIV BERKELEY CENTER FOR PURE
AND APPLIED MATHEMATICS  B N PARLETT APR 87 AAM-372
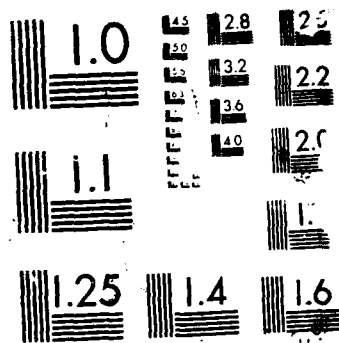
1/1

UNCLASSIFIED    N00044-85-K-0180                        F/G 12/1        NL

# THE CONTRIBUTION OF J. H. WILKINSON

## TO NUMERICAL ANALYSIS[†]

B. N. Parlett[*]

Mathematics Department and
Computer Science Division of EECS Department
University of California
Berkeley, California 94720

## 1. An Outline of His Career

James Hardy Wilkinson died suddenly at his London home on October 5, 1986, at the age of 67. Here is a very brief account of his professional life.

He won an open competition scholarship in mathematics to Trinity College, Cambridge, when he was 16 years old. He won two coveted prizes (the Pemberton and the Mathieson) while he was an undergraduate at Trinity College and graduated with first class honors before he was 20 years old.

He worked as a mathematician for the Ministry of Supply throughout World War II and it was there that he met and married his wife Heather. In 1947 he joined the recently formed group of numerical analysts at the National Physical Laboratory in Bushy Park on the outskirts of London. He was to stay there until his retirement in 1980. Soon after his arrival he began to work with Alan Turing on the design of a digital computer. That work led to

---

the pilot (prototype) machine ACE which executed its first scientific calculations in 1953. Wilkinson designed the multiplication unit for ACE and its successor DEUCE.

One could say that the decade 1947-1957 was the exciting *learning* period in which Wilkinson, and his colleagues at NPL, discovered how automatic computation differed from human computation assisted by desk top calculating machines. By dint *of trying* every method that they could think of and watching the progress of their computations on punched cards, paper tape, or even lights on the control console, these pioneers won an invaluable practical understanding of how algorithms behave when implemented on computers.

Some algorithms that are guaranteed to deliver the solution after a fixed number of primitive arithmetic operations IN EXACT ARITHMETIC can produce, on some problems, completely wrong yet plausible output on a digital computer. That is the fundamental challenge of the branch of numerical analysis that Wilkinson helped to develop.

The period 1958-1973 saw the development, articulation, and dissemination of this understanding of dense matrix computations. It was in 1958 that Wilkinson began giving short courses at the University of Michigan Summer College of Engineering. The notes served as the preliminary versions of his first two books. The lectures themselves introduced his work to an audience broader than the small group of specialists who had been brought together in 1957 by Wallace Givens at Wayne State University, Michigan, for the first of a sequence of workshops that came to be called the Gatlinburg meetings. These conferences are discussed in more detail in the chapter by R.S. Varga. The year 1973 saw the beginning of the NATS project (at Argonne National Laboratory, USA) whose goal was to translate into FORTRAN, and test even further, the ALGOL algorithms collected in the celebrated Handbook of 1971. That book, essentially by Wilkinson and Reinsch, embodied most of what had been

learnt about matrix transformations. There is more on this topic below.

By 1973 Wilkinson had received the most illustrious awards of his career. He was elected to the Royal Society of London in 1969. In 1970 he was awarded both the A. M. Turing award of the Association for Computing Machinery and the John von Neumann award of the Society for Industrial and Applied Mathematics. Both these professional groups are in the USA. It was not until 1977 that he was made an honorary fellow of the (British) Institute for Mathematics and its Applications.

The final period, 1974-1986, may be marked by Wilkinson's promotion to the Council of the Royal Society. Indeed he served as secretary for the physical sciences section for two or three years and these duties absorbed much of his energy. When that obligation was discharged he accepted a professorship in the Computer Science department at Stanford University, California, (1977-1984) but he was only in residence for the Winter quarter and not every year was he able to take up his position. His research now focused on more advanced, but less urgent numerical tasks such as computing the Jordan form, Kronecker's form for matrix pencils, and various condition numbers. During the last four years of his life he was absorbed in the still open problem of how to determine the closest defective matrix to any given square matrix. He also gave much attention to the task of explaining to the wider mathematical community the nature of the subject with which his name is indissolubly linked: roundoff error analysis. We will say more on this expository problem below.

3

## 2. Background

People are awed at the prodigious speed at which the digital computers of the 1980's can execute primitive arithmetic operations; sometimes millions of them per second. Yet this speed is achieved at a price; almost every answer is wrong. When two 14 decimal digit numbers are multiplied together, only the leading 14 digits are retained, the remaining 13 or 14 digits are discarded forever. If such a cavalier attitude to accuracy were to make nonsense of all our calculations then the prodigious speeds would be pointless. Moreover it requires little experience to discover how easily a digital computer can produce meaningless numbers.

Fortunately there are procedures that can survive these arithmetic errors and produce output that has adequate accuracy. Consequently computers can be useful. The difficult task is to discern the robust algorithms. A poor implementation can undermine a sound mathematical procedure and this simple fact has extensive and unpleasant consequences. It suggests that clean, general statements about the properties of numerical methods may not always be possible. Here is an example: will the process of iterative refinement improve the accuracy of the output of a good implementation of Gauss elimination on an ill-conditioned system of linear equations? The answer turns out to depend on whether certain intermediate quantities are computed with extra care. Considerations of this sort make it difficult to present the results of an error analysis and Wilkinson became more and more concerned with this problem.

Before embarking on a list of Wilkinson's contributions, five points must be emphasized.

1)    Only a minority of numerical analysts pay attention to roundoff error. For example, in his influential book *Matrix Iterative Analysis* (Prentice-Hall, 1962), R. S. Varga mentions during the introduction that he will not be considering the effects of roundoff-error.    Virtually every publication concerned with the approximate solution of differential equations invokes exact arithmetic.  The tacit assumption is that the approximation errors are so much greater than the effect of roundoff that the latter may be ignored without risk.

Wilkinson's brand of numerical analysis is perhaps best regarded as an extra layer in the analysis of approximate solutions.  It slips in just above the arithmetic facilities themselves.

2)    The pages that follow give the erroneous impression that Wilkinson single-handedly showed the world how to analyze the effect of roundoff error. Yet this mode of expression is *no worse* than the familiar statement that William of Normandy won the battle of Hastings in 1066.  Wilkinson did receive all the honors and most would agree that he became the leader of the group. Yet he was not working in isolation.  Other people, independently, came to understand how roundoff errors can destroy a computation.  I would like to intrude my personal opinion that had Wilkinson returned to classical analysis at Cambridge in 1947 our present state of understanding of roundoff would not be significantly changed.  F. L. Bauer of Munich could have become the dominant figure, or H. Rutishauser of Zurich.

3)    The production of *The Handbook* was a remarkable achievement.  It testifies to cordial and close cooperation between leading experts in several European countries, the USA, and Australia.   In contrast, consider the application of the simplex algorithm to linear programs and the finite element method to analyze structures.   There it was the habit for engineers with debugged programs to form companies round those programs.  The quest for

profit stifled cooperation for improvement. Wilkinson's friendly yet exacting personality played no small part in the success of the Handbook venture. I am aware of no disharmony among the leading researchers on matrix problems.

4) A digital computer works with a finite set of representable numbers which may be combined using operations $\oplus, \ominus, \otimes, \oslash$ that mimic the familiar $+$, $-$, $\times$, $\div$. Unfortunately, some basic properties of the rational number field fail to hold for the computer's system. For example, the associative law fails for both addition and multiplication. Nevertheless, there is some algebraic structure left and it seemed quite likely during the 1950's that rigorous error analysis would have to be carried out in this unattractive setting. Indeed there have appeared a number of ponderous tomes that do manage to abstract the computer's numbers into a formal structure.

Perhaps Wilkinson's greatest achievement was to deflect analysis of algorithms from that morass into a place where insight and simplicity can survive. He makes no use of the pseudo-operators preferring to work with the exact relations satisfied by the computed quantities.

5) In contrast to most mathematicians and despite over 100 published papers, Wilkinson's contribution to numerical analysis is contained in the three books of which he is an author.

## 3. Roundoff Error Analysis

Wilkinson is honored for achieving a very satisfactory understanding of the effect of rounding errors during the execution of procedures that are used for solving matrix problems and finding zeros of polynomials. He managed to share his grasp of the subject with others by making error analysis intelligible, in particular by systematic use of the "backward" or inverse point of view. This approach asks whether there is a tiny perturba-

tion of the data such that execution of the algorithm in exact arithmetic using the perturbed data would terminate with the actual computed output derived from the original data. This approach appears in [Turing, 1948].

Wilkinson did not invent backward error analysis nor did he refrain from using the natural (or forward) error analysis when convenient. Although his name is not associated with any particular method he performed rigorous analyses of almost every method that was under discussion and trial. This work led him to become one of the leaders of an activity known as mathematical software production. The collection of Algol procedures contained in *The Handbook* (see reference list) is a seminal contribution to this branch of computer science.

Most of what follows is amplification of the preceding paragraphs. If the reader is impatient for a theorem or delicate inequality, the following quotation from *Modern Error Analysis* (1971) may engender a little forbearance. This is from the published version of his von Neumann lecture.

"There is still a tendency to attach too much importance to the precise error bounds obtained by an a *priori* error analysis. In my opinion, the bound itself is usually the least important part of it. The main object of such an analysis is to expose the potential instabilities, if any, of an algorithm so that, hopefully, from the insight thus obtained one might be led to improved algorithms. Usually the bound itself is weaker than it might have been because of the necessity of restricting the mass of detail to a reasonable level and because of the limitations imposed by expressing the errors in terms of matrix norms. A *priori* bounds are not, in general, quantities that should be used in practice. Practical error bounds should usually be determined by some form of a *posteriori* error analysis, since this takes full advantage of the statistical distribution of rounding errors and of any special features, such as sparseness, of the matrix."

We would add that there is as yet no satisfactory format for presenting an error analysis so that its message can be summarized succinctly. Could we say that the analysis is the message?

## 4. The Linear Equations Problem

Given an $n \times n$ real invertible matrix $A$ and $b \in \mathbb{R}^n$ the task is to compute $x = A^{-1}b$. The familiar process known as Gaussian elimination lends itself to implementation on automatic digital computers. It is also well known that Gaussian elimination is one way to factor $A$ into the product $LU$ where $L$ is lower triangular and $U$ is upper triangular. Once $L$ and $U$ are known the solution $x$ is obtained by solving two triangular systems:
$Lc = b, \quad Ux = c$.

In 1943, Hotelling published an analysis showing that the error in a computed inverse $X$ might well grow like $4^{n-1}$ where $n$ is the order of $A$. Alan Turing was making similar analyses informally in England. The fear spread that Gaussian elimination was probably unstable in the face of roundoff error. The search was on for alternative algorithms.

In 1947 Goldstine and von Neumann, in a formidable 80-page paper published in the *Bulletin of the American Mathematical Society*, corrected this false impression to some extent. Some scholars have chosen the appearance of this paper as the birthday of modern numerical analysis. Among other things, this paper showed how the systematic use of vector and matrix norms could enhance error analysis. However it had the unfortunate side effect of suggesting that only people of the calibre of von Neumann and Goldstein were capable of completing error analyses and, even worse, that the production of such work was very boring. Their principal result was that, if $A$ is symmetric, positive definite, then the computed inverse $X$ satisfies

8

$$\|AX - I\| \leqslant (14.2)n^2 \varepsilon \, cond(A)$$

where $cond(A) = \|A\| \cdot \|A^{-1}\|$, and $\|\cdot\|$ is the spec- tral norm and $\varepsilon$ denotes the roundoff unit of the computer. Only if A is too close to singular will the algorithm fail and yield no X at all, but that is as it should be. The joy of this result was getting a polynomial in n, the pain was obtaining 14.2, a number that reflects little more than the exigencies of the analysis. Some nice use of "backward" error analysis occurs in the paper but it is incidental. There was good reason for this attitude.

A backward error analysis is not guaranteed to succeed. Indeed no one, to this day, has shown that a properly computed inverse X is guaranteed to be the inverse of some matrix *close* to A, i.e.,

$$X = (A + E)^{-1} \text{ and } \|E\|/\|A\| \text{ is small.}$$

Indeed, it is likely that no such result holds in full generality. What is true is that each column of X is the corresponding column of the inverse of a matrix very close to A. Unfortunately, it is a different matrix for each column.

The success of their analysis of the positive definite case prompted von Neumann and Goldstein to recommend the use of the normal equations for solving Ax = b for general A, i.e., $x = (A^T A)^{-1} A^T b$. However, that was bad advice for several reasons.

The fact is that careful Gaussian elimination, if it does not break down, produces computed solutions z with tiny residuals. It was practical experience in solving systems of equations using desk-top calculators (with n as large as 18)! that persuaded Wilkinson and his colleagues (L.Fox and E. T. Goodwin) that Gaussian elimination does give excellent results even when A is far from being symmetric let alone positive definite. In his first book *Rounding Errors in Algebraic Processes*, published in 1963, we find for the first time a clear statement of the situation (see p.108). The computed

9

solution $z$ satisfies

$$(A + K)z = b \quad .$$

If inner products are accumulated in double precision before the final rounding then

$$\|K\|_\infty \leq g\varepsilon(2.005 \, n^2 + n^3 + \tfrac{1}{4}\varepsilon n^4)\|A\|_\infty$$

where $g$ is the element growth factor, namely, the ratio of the largest intermediate value generated in the process to a maximal element of A. The corresponding bound on the residual is

$$\|b - Az\|_\infty \leq g\varepsilon(2.005 \, n^2 + n^3)\|z\|_\infty$$

provided that $\varepsilon n \ll 1$. The important quantity $g$ is easily monitored during execution of the algorithm. In his celebrated 1961 paper on matrix inversion, Wilkinson obtains an *a priori* bound on $g$ when A is equilibrated and the "complete" pivoting strategy is employed. This is a clever piece of analysis and yields:

$$g^2 = g(n)^2 < n(2^1 3^{\frac{1}{2}} \cdots n^{\frac{1}{n-1}}) \quad ,$$

a *slowly growing* function of n. Being a man of intellectual integrity Wilkinson hastens to show that the bound cannot be sharp and indeed is *not* realistic at all. For certain Hadamard matrices $g(n) = n$, but apart from these cases Wilkinson reports that he has *never* encountered a value of $g$ exceeding 8 despite intensive monitoring of the programs in use at NPL.

At this point we wish to emphasize that all the results quoted so far do a disservice both to Wilkinson and the topic of error analysis. Neither the powers of n that appear in the inequalities quoted above nor the coefficients in front of those powers convey genuine information about the process under analysis! It could be argued that the residual bound $\|b-Az\| <$

10

$g\varepsilon n^3 \|z\|$ is very weak indeed. Wilkinson's contribution cannot be conveyed by quoting such theorems. His achievements in regard to Gaussian elimination was to show the following:

- The effect of roundoff errors is not difficult to analyze. Indeed, the analysis is now presented in undergraduate courses.

- If the element growth factor $g$ is small (say, $g < 5$) then the computed solution will have a residual norm scarcely larger than that belonging to the representable vector closest to $A^{-1}b$.

- When $A$ is ill-conditioned, i.e., $\|A\| \; \|A^{-1}\| \sqrt{\varepsilon} > 1$ then $g$ is very likely to be 1 if a reasonable pivoting strategy is used. In fact, for many ill-conditioned matrices the complete pivoting strategy produces factors $L$ and $U$ with elements that diminish rapidly as the algorithm proceeds.

- The technique known as iterative refinement may be employed to obtain an accurate solution provided that the system is not too ill-conditioned for the precision of the arithmetic operations. Moreover if the iteration converges slowly then the coefficient matrix $A$ must be ill-conditioned.

- The partial pivoting strategy cannot guarantee that $g$ will be small. There exist matrices for which $g = 2^{n-2}$.


The following very specific result of the 1961 paper is, to me, more interesting and more informative than all its theorems. The Hilbert matrix was a favorite test example in the 1940's and 1950's,

$$H = (h_{ij}) \; , \qquad h_{ij} = (i+j-1)^{-1}$$

$H_n$ denotes the leading principal $n \times n$ submatrix of H. Formulae are known for $H_n^{-1}$. Wilkinson showed that when Gaussian elimination was used to invent $H_5$ on a binary machine then the act of rounding the fractions 1/3, 1/5, 1/6, 1/7,

11

1/9 to the closest representable numbers caused *more* deviation in the computed inverse than *all* the rounding errors that occur in the rest of the computation. That computation involves more than 100 multiplications and 100 additions.

Despite several significant insights, the celebrated 1961 paper still does not make clear just how stable Gaussian elimination is for solving $Ax = b$. The contrast between this paper and the 1963 book is instructive. The paper follows the lead of von Neumann and Goldstine and concentrates exclusively on the problem of matrix inversion. Not only are the error bounds rather large but backward error analysis fails. However the problem of matrix inversion is not very important. The overwhelming demand is for solving systems of equations and here the backward analysis is simple and very satisfactory. The computed solution $z$ satisfies some equation $(A + K)z = b$ and the insight comes in seeing how $K$ depends on $L$, $U$ and other quantities. The insight vanishes when norms are taken. Too much information is discarded.

## 5. The Eigenvalue Problem

Nearly three quarters of Wilkinson's publication list is devoted to this subject. No specific method bears his name yet every available method was analyzed by him and most of the published implementations of the better techniques owe something to his careful scrutiny.

The eigenvalue problem comprises many subproblems. The primary distinction is between symmetric matrices and the rest. For both classes the eigenvectors may or may not be needed. It is easy to describe Wilkinson's contribution to this topic. It is his magnum opus, *The Algebraic Eigenvalue*

*Problem* (Oxford University Press, 1965). However that gives only half the picture. That book gave the understanding needed to produce the eigenvalue programs that appeared in the Handbook (*Handbook for Automatic Computation*, vol. II, *Linear Algebra*, edited, Springer-Verlag, 1971). The latter was edited jointly with the gifted but self-effacing Dr. Christian Reinsch. The Handbook gave rise to the collection of FORTRAN programs called EISPACK which first appeared in 1974. The later version of these routines (1977) are available in virtually every scientific computer center in the world. It is pleasant to report that this useful product was achieved by the willing cooperation of many experts. To some extent this happy outcome is due to Wilkinson's generous and agreeable personality for he was certainly the leader of the group.

At a more technical level we now discuss some of his "results". Journal articles are given in the brief reference list.

● His study of polynomials, and the sensitivity of their zeros to changes in their coefficients, helped to stop the quest for the characteristic polynomial as a means of computing eigenvalues. See *Rounding Errors in Algebraic Processes.*

● In 1954, W. Givens explicitly used backward error analysis to demonstrate the extreme accuracy of the Sturm sequence technique for locating specified eigenvalues of symmetric tridiagonal matrices. His analysis was for fixed point arith- metic and was never published. Wilkinson showed that the result still holds for standard floating point arithmetic and, contrary to popular wisdom, that backward error analysis of most algorithms is easier to perform for floating point arithmetic. Even more interesting was his demonstration that Givens Sturm sequence algorithm could be disastrous for computing eigenvectors while simultaneously being superb for locating

13

eigenvalues. The point is worth emphasizing. Given an appropriate eigenvalue that is correct to working precision the eigenvector recurrence can sometimes produce an approximate eigenvector that is orthogonal to the true direction to working accuracy yet the signs of the computed components are correct. The contribution here was a well chosen class of examples.

● Wilkinson showed that the backward error analysis of *any* method employing a sequence of orthogonal similarity transformations can be made clear and simple. In particular the final matrix is similar to a small perturbation of the original matrix. This perturbation is essentially the sum of the local errors at each step: there is no propagated error.

An important consequence of this analysis is the following. Let C denote the equivalence class of matrices orthogonally similar to the original matrix. For the computation of eigenvalues it does not matter if roundoff errors cause the computed sequence to depart violently from the exactly computed sequence *provided that* the computed sequence lies close to C. A naive forward analysis can miss vital correlations between computed quantities.

Indeed a number of efficient, stable algorithms do regularly produce intermediate quantities that differ significantly from their exact counterparts. Nevertheless eigenvalues are preserved to within working accuracy. The QR algorithm is an example of this phenomenon.

● Although it was invented in 1959/60, the QR algorithm of J.G.F.Francis did not achieve universal acceptance until about 1965. It provides an ideal way to diagonalize a symmetric tridiagonal matrix since it produces a sequence of symmetric tridiagonal matrices that converge to diagonal form. However the QR algorithm requires a strategy for choosing shifts. Let

$$T_n = \begin{bmatrix} \alpha_1 & \beta_1 & & & & \\ \beta_1 & \alpha_2 & \beta_2 & & & \\ & \beta_2 & \cdot & \cdot & & \\ & & \cdot & \cdot & \beta_{n-1} & \\ & & & \beta_{n-1} & \alpha_n \end{bmatrix}$$

be such a matrix with $\beta_i > 0$, $i = 1,...,n-1$. Wilkinson's shift $\omega$ is defined to be the eigenvalue of $\begin{pmatrix} \alpha_{n-1} & \beta_{n-1} \\ \beta_{n-1} & \alpha_n \end{pmatrix}$ that is closer to $\alpha_n$. It is the favorite strategy (rather than choosing $\alpha_n$) but when it was first introduced there was no proof that it would always lead to convergence. Convergence here means that $\beta_{n-1} \to 0$ as the algorithm is continued without limit. The Rayleigh quotient shift $\alpha_n$ causes the $\beta_{n-1}$ to be monotone decreasing but the limit need not vanish. Wilkinson's shift sacrifices the monotonicity and gains convergence.

In a tour-de-force in 1971 Wilkinson proved that, with his strategy, convergence is assured (in exact arithmetic) and is usually cubic. A tricky argument showed that the product $\beta_{n-2}\beta_{n-1}$ is monotone decreasing to zero though initially the rate could be very slow.

This was not the last word however. In 1979 Parlett and Hoffman discovered an elementary proof that $\beta_{n-1}^2 \beta_{n-2}$ decreases geometrically at each step by a factor at most $1/\sqrt{2}$. Convergence of $\beta_{n-1}$ follows readily.

● In 1976, Wilkinson and Golub published a long article on the Jordan canonical form. They discussed its discontinuous dependence on the matrix elements in the defective case. They showed how to go about computing robust bases for the associated cyclic subspaces (the numerical analyst's Jordan chains of principal vectors), and they also explained the limitations of this form in practical calculations.

A natural extension of this research was to the computation of the Kronecker form of a pair of matrices (A,B). This form arises in the study of

systems of differential equations with constant coefficients:

$$B\dot{u} = Au \ , \qquad u(0) \in \mathbb{R}^n \ \text{ is given } .$$

● In the last decade of his life Wilkinson's attention was more and more attracted to the difficult and still open problem of determining, *for any given* A, the closest defective matrix B.

## 6. The Zeros of Polynomials

Until near the end of the 1950's the computation of the zeros of polynomials was regarded as an important activity in scientific computation. So it is not surprising that a significant part of Wilkinson's work of this period was devoted to this task. His contribution is consolidated in Chapter 2 of *Rounding Errors in Algebraic Processes*. Thanks in part to his discoveries, polynomials no longer attract much attention. It was the advent of digital computers that drove people to think in detail about general polynomials of large degree; 20 or 100 or 1000.

Since isolated zeros are analytic functions of the coefficients one may consider the derivative of any isolated zero with respect to each coefficient. As the degree rises these derivatives can hardly avoid becoming huge. The presence of such an ill-conditioned zero can make it difficult to compute comparatively well conditioned zeros.

By use of well chosen examples Wilkinson brought these facts home to numerical analysts. Of considerable personal interest is the fact that Wilkinson was led to an explicit appreciation of the importance of backward error analysis when he investigated the reliability of Horner's method (also known as nested multiplication) for *evaluating* a polynomial. He realized that,

16

with floating point arithmetic, the output of Horner's recurrence is, in all cases, the exact value of a polynomial each of whose coefficients is a tiny relative perturbation of the original one. The relative change in the coefficient of $x^r$ is less than

$$(1.01)(r+1)2\varepsilon \quad ,$$

where $\varepsilon$ is the roundoff unit. In the majority of cases the inherent uncertainty in each coefficient will exceed the worst case error given above. In this way a fearsome error analysis melts away into classical perturbation theory.

One of Wilkinson's final works, "The Perfidious Polynomial" (Chapter I in *Studies in Numerical Analysis*, G. H. Golub, editor, Math. Assoc. Amer., vol. 24, 1984) sums up his experience with polynomials in a way that is designed for readers outside numerical analysis. This pellucid essay was awarded the Chauvenet prize for mathematical exposition. Unfortunately Wilkinson died before he could receive it.

REFERENCES

Books by JHW

*Rounding Errors in Algebraic Processes*, Notes on Applied Science No. 32, National Physical Laboratory, England, HMSO, 1963. Also published in 1964 by Prentice-Hall, Inc.

*The Algebraic Eigenvalue Problem*, Oxford University Press, 1965.

*Handbook for Automatic Computation*, vol. II: *Linear Algebra*, (edited with C. Reinsch), Springer-Verlag, 1971.

Selected Articles by JHW

"Error Analysis of Direct Methods of Matrix Inversion", *Journal of the Association for Computational Machinery*, vol. 8 (1961), 281-329.

"Error Analysis of Eigenvalue Techniques Based on Orthogonal Transformations", *Journal of SIAM*, vol. 10 (1962), 162-195.

"Error Analysis of Eigenvalue Techniques Based on Orthogonal Transformations", *Journal of SIAM*, vol. 10 (1962), 162–195.

"Global convergence of tridiagonal QR algorithm with origin shifts", *Linear Algebra and Its Applications*, vol. 1 (1968), 409.

(with G. H. Golub), "Ill-conditioned eigensystems and the computation of the Jordan canonical form", *SIAM Review*, vol. 18 (1976), 578.

"Kronecker's canonical form and the QZ algorithm", *Linear Algebra and Its Applications*, vol. 28 (1979), 295–305.

"On neighbouring matrices with quadratic elementary divisors", *Numerical Mathematics*,
vol. 44 (1984), 1–21.


## Articles Not by JHW

H. H. Goldstine and J. von Neumann, "Numerical Inventing of Matrices of High Order", *Bulletin of the American Mathematical Society*, vol. 53 (1947), 1021–1099.

W. Hoffman and B. N. Parlett, "A New Proof of Global Convergence for the Tridiagonal QL Algorithm", *SIAM Journal of Numerical Analysis*, vol. 15 (1978), 929–937.

H. Hotelling, "Some new methods in matrix calculations", *Annals of Mathematical Statistics*, vol. 14 (1943), 1–34.

A. M. Turing, "Rounding-off Errors in Matrix Processes," *Quart. J. Mech. Applied Mathematics*, vol. 1 (1948), 287–308.

END

FILMED

5-89

DTIC